



PGI/EP200 4 / 0 0 3 4 6 5



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ


REC'D 11 JUN 2004
WIPO PC

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated 11 May 2004

**PRIORITY
DOCUMENT**

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

Patents Form 1/77

Patents Act 1977
(Rule 16)THE PATENT OFFICE
JG

1 APR 2003

RECEIVED BY FAX

The
Patent
Office

PCT/EP200 4 / 0 0 3 4 6 5

01APR03 E796570-1 D00393
PG1/7700 0.00-0307384.8**Request for grant of a patent***(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)*

The Patent Office

Cardiff Road
Newport
Gwent NP23 5RH**1. Your reference**

P/64000.GB

2. Patent application number*(The Patent Office will fill in this part)*

0307384.8

3. Full name, address and postcode of the or of each applicant (underline all surnames)Telnic Limited
8 Wilfred Street
London SW1E 6PL
United KingdomPatents ADP number *(if you know it)*

If the applicant is a corporate body, give the country/state of its incorporation

8096794001
United Kingdom**4. Title of the invention**

A system for enhancing a publishing data system such as a Domain Name Server

5. Name of your agent (if you have one)

N G McGowan

"Address for service" in the United Kingdom to which all correspondence should be sent *(Including the postcode)*Intellectual Property Department
Siemens Shared Services
Siemens House
Oldbury, Bracknell
Berkshire, RG12 8FZ
United KingdomPatents ADP number *(if you know it)*

7761000002

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
*(if you know it)*Date of filing
*(day / month / year)***7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application**

Number of earlier application

Date of filing
*(day / month / year)***8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer Yes/No)**

Yes

- a) any applicant named in part 3 is not an inventor, or
 b) there is an inventor who is not named as an applicant, or
 c) any named applicant is a corporate body.
 See note (a))

Patents Form 1/77

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form -

Description 7

Claim(s) -

Abstract -

Drawing(s) 1 only

10. If you are also filing any of the following, state how many against each item.

Priority documents -

Translations of priority documents -

Statement of inventorship and right to grant of a patent (Patents Form 7/77) -

Request for preliminary examination and search (Patents Form 9/77) -

Request for substantive examination (Patents Form 10/77) -

Any other documents (please specify) -

11.

I/We request the grant of a patent on the basis of this application.

Signature

N G McGowan

Date 1st April 03

N G McGowan

12. Name and daytime telephone number of person to contact in the United Kingdom

N G McGowan - 01344 396808

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

Patents Form 1/77

A system for enhancing a publishing data system such as a Domain Name Server

1. Abstract

A system to enhance the value of items stored within publicly available data areas, by supporting association of data areas through providing selectively available supplementary information that can be used to give different structural forms for the publicly available data, provision of discrete sets of supplementary data being dependent on the distinct populations of users requesting it.

2. System description

2.1 Prior Art

The Internet is a data network that has become more and more pervasive over recent years. The Internet utilizes the service of Domain Name Servers (DNS Servers) for uniquely locating resources on the Network. DNS servers on the Internet link names (e.g. yahoo.com) with Internet Protocol (IP) addresses (e.g. 192.168.1.1) for establishing data communication.

DNS servers can also be populated with Name Authority Pointer (NAPTR) records that store various Contact Addresses from several Networks (Public Switched Telephony Networks, Web Addresses, Email Addresses, etc). A DNS server can be viewed as a public data store and publisher system that publishes information items to all requesting clients regardless of their identity. Information in this data store can be resolved by a DNS client.

DNS Servers can be viewed as a globally distributed, loosely coherent, scalable, reliable, dynamic database. This database has a fixed structure corresponding to an inverted tree with the root node at the top. Each node of the tree has a label (the root node having the null label). A domain name is the sequence of labels from a node to the root, separated by dots, read left to right, one domain being a sub domain of another if its domain name ends in the other's domain name (yahoo.com is a sub domain of com). Name servers cooperate to publish the data of the name space. In the case of a DNS system the hierarchy and the organization of the data store is fixed and organized according to a Tree structure. The only way of retrieving data published through the DNS is by walking down the name space Tree. The Tree structure of the DNS fixes the hierarchy of domain names in the name space.

2.2 System components

The system has five sub-systems:

Sub-system 1: a public data store and publisher sub-system that stores and publishes information items to all requesting clients regardless of their identity, this data store being partitioned into discrete areas each of which is associated with a

unique area identity, and which can return data stored within this area on request from any user who passes the area identity of interest to the data store.

Sub-system 2: a client sub-system, by which an end user may query the public data store for any information it contains by passing a data area identifier, but in addition may request supplementary information from a separate directory sub-system, passing the identifier of the data area in which they are interested and their identity to that directory so that it may select the supplementary information returned based on their interest and their identity.

Sub-system 3: a directory sub-system that stores structural information holding relations and associations between the public data areas stored and published from the public data store, and presents this information based both on the (public) data area of interest and on the identity of the requesting party. There may be several independent sets of structural information stored for a given (publicly available) data item. Note that this directory does not hold copies of data items, but instead holds only references to the data areas (the area identifiers) and the relations between those data areas. This directory sub-system can be seen as a system enabling navigation in the Public Data Store through the relationships it creates between the public data areas.

Sub-system 4: an editing sub-system by which relations between public data areas may be captured and stored in the directory sub-system, and by which classification of the availability of this data set may be controlled based on the identity of requesting users, and by which authorized users may store data items to be published in an area of the public data store; authorization being granted exclusively to a given user for a particular area of the data store.

Sub-system 5: a query engine that can return a reference to a public data area matching a passed search term.

Alternative 1: the incorporation of an additional registry controlled sub-structure (e.g. a sub-domain in the case of a DNS system) can permit enhanced management and query of the public and/or directory data.

Alternative 2: an integrated management level used by the end users, whereby they can manage their specific data according to their own logic.

Alternative 3: as a variation on the above, the directory sub-system may in addition store and publish data areas (with their constituent data items) in a similar form to those published from the public data store sub-system; thus not only does the directory hold structural information but also individual data areas that can act as new nodes in the structure that are only available from within the directory.

Alternative 4: the inclusion of an End-User system to identify, differentiate, and resolve between, contact information (telephone numbers, email addresses, etc) and content information (web site addresses – personal or corporate, web based application addresses, etc).

2.3 System Diagram

This is shown in the accompanying drawing.

3. Example embodiment:

3.1 Definitions

3.1.1 Subscriber Name

The Subscriber Name is a domain name that is within the hierarchy of a name space together with hierarchical or contact information associated with that domain.

3.1.2 Registrant

A person that purchases the Subscriber Name through the Client Subsystem.

3.1.3 TLD Registry

The entity that has authority over the name space by managing the top of the DNS Tree.

3.1.4 Domain Name System (DNS)

A service provided by distributed directory servers by which resources can be associated with a node within a single, global hierarchical name space, and can be queried and returned using a defined protocol. The overall definitions for the DNS are specified in RFC1034 and RFC1035.

3.1.5 Domain

The delegated Domain Name Service domain associated with a Subscriber Name.

3.1.6 Subscriber Contact Information

Associated with a Subscriber Name is a set of records each of which holds some information on the communications contacts defined for the assignee of that domain. Where published in the Domain Name System, this will be stored within Name Authority Pointer (NAPTR) resource records. In addition, this Contact Information may be published within other services, such as Web-based "Whois" Servers, Web-based Search Engines, or in Directory Servers.

3.1.7 Fully Qualified Domain

A Fully Qualified Domain is the leaf of a domain space based on the top level domain. Only a Fully Qualified Domain has contact information associated with it that can be returned via the Domain Name Service.



3.1.8 Directory

A Hierarchical database set (potentially distributed) that stores Subscriber Names and their associated contact information items. This database stores records holding detailed content for a delegated Subscriber Name Domain (or set of such domains). This Server publishes these records using the DNS protocol. This Directory holds references to distributed Servers that in turn store records with the content for a Subscriber Name Domain delegated by the TLD Registry.

3.1.9 DNS Tree

A hierarchical distributed database that stores Subscriber Contact Information items that are published to clients everywhere using the DNS protocol. Entries in the DNS Tree use standard resource record types (DNS RR records).

3.1.10 Directory Tree

This Structure may be defined as a set of relationships between nodes in the Directory repository.

3.1.11 Node

A node has two sets of relations; 'parent' relations that refer to one set of nodes, and 'child' relations that refer to another set of nodes. In a tree structure, exactly one node has an empty parent relation set. All other nodes have one parent relation. A node may have zero or more child relations. These rules form the definition of a tree structure. When discussing tree structures, the node with no parent is often referred to as the "root" of the tree, whilst those nodes that have no child relations are known as "leaves" of the tree. Nodes that have both a parent relation and child relations are known as "branches" of the tree.

3.2 Description

The public data store and publisher sub-system is realised using a device that implements the standard Domain Name Service (DNS, as specified in RFC1034 and RFC1035), and in addition stores the names of the person responsible for each data area. In this case the discrete data areas it holds are domains, and the identifier for these domains are domain names. The data items stored within the areas are DNS Resource Records.

DNS is designed as a distributed hierarchical data store; either the "data store and publisher sub-system" device holds the data items directly within its data store, or it holds a reference to an external (subsidiary) DNS server that supports that data area. These subsidiary servers may in turn support sub-areas within the data area they hold. Such sub-areas are identified by identifiers that are within the context of the identifier for the area that "contains" them. In DNS, this is done by prepending a further label to the main identifier, with a separator between the label and the main identifier. As a hierarchical system, the subsidiary DNS servers may hold the data items associated with such sub-areas themselves, or hold references to another DNS

server in which the enclosed sub-area and its data items are stored, as specified in RFC1591. Note: See RFC1034 for definitions of these terms.

When a new domain is created within the main DNS device in this system, an automatic notification is sent to the query engine and to the directory listing the domain name with which the domain is associated, along with the name of the person responsible for that domain's creation; thus by this notification they can build a complete list of the domain names known to this DNS device, and the names of the people responsible for those domains. Note, however, that they are not informed of the status or values of any data items that may be stored within the domains so identified.

The query engine is realised using a device that accepts a textual value as a search term, applies this value to its database of names of people responsible for domain names, and returns the domain name or names for which a person whose name matches the passed search term is responsible, or a list of domain names associated with persons whose name forms "closest matches" to the search term, using a standard affinity pattern matching algorithm. The query engine has a complete list of all domain names and their responsible person names known to the DNS sever, as an indication of the creation of all new domains (and the domain name with which they are associated) is sent from the main DNS device to the query engine.

The directory sub-system is realized using a device that stores relation data sets that are associated with individual domain names. In one example, these consist of a set of references to other domain names that form a tree structure based on the domain name in question, together with an access control policy that identifies the class of querying users to whom this structural data should be returned. The directory also receives notification on the creation of new domains within the main DNS device, so its internal data store includes all of the domain names associated with domains held in the main DNS device.

The directory holds references to not only the domain names known in the main DNS device but also to subsidiary domain names that have been introduced by authorized users; these are sub-domain names within the context of the domain name notified on creation from the main DNS device.

Thus the directory is informed of other domain names that exist outside of the DNS device.

Thus the structural data the directory returns may include reference not only to those domain names and responsible person names stored in the main DNS device, but also to those held in external subsidiary DNS devices.

Finally, the structures may include intermediary nodes in the tree structure; such nodes need not exist outside the directory itself and exist only as constructs of relational data sets.

It is novel to use automatic notification from a trusted source - i.e. the Registry - to pre-populate a directory with domain names and the identity of the person responsible for their creation. This allows the directory to (i) ensure that relation data

is based on the domains that exist in the external DNS system, (ii) control the ability of external users to create relations based on those domains - i.e. only the person responsible for the creation of that domain is allowed to create relations based on it.

The relationship between a Registry and a Directory Service Provider using this automatic notification improves the efficiency of the directory considerably, whilst minimizing the privileged data that flows through the notifications. When a new domain has been created the information on that domain and on the person responsible for its creation is only known to the Registry. The Directory Service Provider gets this information through this certified automatic notification process.

The relationship between a Registry and a Directory Service Provider using this automatic notification process simplifies the directory's subsequent task of capturing data on relations between these names, and additionally other names that are added as part of this process, in that the directory receives these automatic notifications from a trusted source; thus the names that can appear at the root of these relational structures are already in place, along with the identities of the persons responsible for these areas. Thus the directory can apply controls over the person who can construct the relational data that the directory will subsequently present based on this notified data.

This supplementary relational data can be used by querying users to relate names to one another. As there can be more than one set of relational data for a given domain, the data returned can be based on the identity of the querying user.

Having the identity of the person responsible for creation of the domain that is to be the base of the relational data means that the directory can easily ensure that only that person has a right to construct supplementary relation data concerning the domain, and so the directory can enforce this with information it already has, through this automatic notification process.

It is novel to use automatic notification from a trusted source to populate a search-engine with names of the persons responsible for having a domain created. This allows a querying user to search based on the name of that responsible person, and for the search engine to return the domains with which they are associated, or to use affinity-based searches to return the domains associated with people whose names are the "nearest matches" for the name passed by the user as a search term. The relationship between the Registry and the search engine provider improves the efficiency of the search engine sub-system; it can be considerably automated using these notifications, whilst it would require at least polling the Registry to find out if new domains had been created and if so at whose request if these notifications were not done.

Additionally, allowing the directory to pass notifications of other domains that it has been informed of by its authorised users (along with their identity and the access controls they wish applied) means that the relationship between the directory and the search engine allows the search engine to be completely automated. No external interaction to control its data from others is required.

It is novel that the directory and the DNS device, in conjunction, form a hybrid system in which the directory stores only relation data associating names, whilst the data items held within data areas that are identified by these names are stored externally to the directory in a separate distributed hierarchical data store; the DNS. The directory is informed of sub-domains that would not be known to the Registry's DNS device by the persons responsible for the domains and their sub-domains. This allows those responsible persons exclusive control over this relation data, and allows the directory to publish this data selectively based on the class of user requesting this supplementary information and on the access control policies specified by the domains' responsible persons.

The configuration of this hybrid system allows the domains and their contained items to be held on a distributed data store that provides one set of data, whilst storing supplementary sets of relational information within a directory that can select which set to return based on the querying user's identity as well as the domain in which they are interested.

This has the benefit of maintaining control for publication of the contained items within the distributed data store; the directory does not store these items but only references to the domain name identifiers. It also ensures that there is only one copy of the data items, whilst allowing different "views" on the relationships between the domains that contain the items based on a querying user's identity.

Having only one copy of the data items removes a problem of synchronization between different systems holding copies. However, information on the hierarchy that relates different domains (that is not normally available to the public through the DNS system) is stored separately in the directory where it can be provided with controlled access, with different information being provided depending on the identity of the user asking for it.

By reflecting the hierarchy of domain names, the person responsible for the enclosing domain name is identifiable (as this identity is passed when the domain for which they are responsible is created). Thus control over the rights to relate domains and their sub-domains is controlled; it remains with that responsible person.

2.3 System diagram

